

17 Sep 2018 | 20:09 GMT

## David Patterson Says It's Time for New Computer Architectures and Software Languages

Moore's Law is over, ushering in a golden age for computer architecture, says RISC pioneer

By **Tekla S. Perry** (/author/perry-tekla-s)



Photo: Peg Skorpinski/UC Berkeley  
David Patterson

David Patterson—University of California professor, Google engineer, and RISC pioneer (<http://news.berkeley.edu/2018/03/21/david-patterson-pioneer-of-modern-computer-architecture-receives-turing-award/>)—says there's no better time than now to be a computer architect.

That's because Moore's Law really is over, he says: "We are now a factor of 15 behind where we should be if Moore's Law were still operative. We are in the post-Moore's Law era."

This means, Patterson told engineers attending the 2018 @Scale Conference (<https://atscaleconference.com/events/the-2018-scale-conference/>) held in San Jose last week, that "we're at the end of the performance scaling that we are used to. When performance doubled every 18 months, people would throw out their desktop computers that were working fine because a friend's new computer was so much faster."

But last year, he said, "single program performance only grew 3 percent, so it's doubling every 20 years. If you are just sitting there waiting for chips to get faster, you are going to have to wait a long time."

**"There are Turing Awards waiting to be picked up if people would just work on these things."**

—David Patterson

For a computer architect like Patterson, this is actually good news. It's also good news for innovative software engineers, he pointed out. "Revolutionary new hardware architectures and new software languages, tailored to dealing with specific kinds of computing problems, are just waiting to be developed," he said. "There are Turing Awards waiting to be picked up if people would just work on these things."

As an example on the software side, Patterson indicated that rewriting Python into C gets you a 50x speedup in performance. Add in various optimization techniques and the speedup increases dramatically. It wouldn't be too much of a stretch, he indicated, "to make an improvement of a factor of 1,000 in Python."

On the hardware front, Patterson thinks domain-specific architectures just run better, saying, "It's not magic—there are just things we can do." For example, applications don't all require that computing be done at the same level of accuracy. For some, he said, you could use lower-precision floating-point arithmetic instead of the commonly used [IEEE 754](https://en.wikipedia.org/wiki/IEEE_754) ([https://en.wikipedia.org/wiki/IEEE\\_754](https://en.wikipedia.org/wiki/IEEE_754)) standard.

The biggest area of opportunity right now for applying such new architectures and languages is machine learning, Patterson said. "If you are a hardware person," he said, "you want friends who desperately need more computers." And machine learning is "ravenous for computing, which we just love."

Today, he said, there's a vigorous debate surrounding which type of computer architecture is best for machine learning, with many companies placing their bets. Google has its [Tensor Processing Unit \(TPU\)](https://cloud.google.com/blog/products/gcp/an-in-depth-look-at-googles-first-tensor-processing-unit-tpu) (<https://cloud.google.com/blog/products/gcp/an-in-depth-look-at-googles-first-tensor-processing-unit-tpu>), with one core per chip and software-controlled memory instead of caches; [Nvidia](https://www.nvidia.com/en-us/) (<https://www.nvidia.com/en-us/>)'s GPU has 80-plus cores; and Microsoft is taking an FPGA approach.

And [Intel](http://www.intel.com/) (<http://www.intel.com/>), he said, "is trying to make all the bets," marketing traditional CPUs for machine learning, [purchasing Altera](http://fortune.com/2015/08/27/why-intel-altera/) (<http://fortune.com/2015/08/27/why-intel-altera/>) (the company that provides FPGAs to Microsoft), and [buying Nervana](https://www.engadget.com/2017/10/17/intel-ai-deep-learning-nervana-npp/) (<https://www.engadget.com/2017/10/17/intel-ai-deep-learning-nervana-npp/>), with its specialized neural-network processor (similar in approach to Google's TPU).

Along with these major companies offering different architectures for machine learning, Patterson says there are at least 45 hardware startups tackling the problem. Ultimately, he said, the market will decide.

"This," he says, "is a golden age for computer architecture."

## The View From the Valley Newsletter

The people, places and peculiarities of Silicon Valley.

### About View From the Valley blog

*IEEE Spectrum's* blog featuring the people, places, and passions of the world of technologists in Silicon Valley and its environs.

**Tekla Perry** (</author/perry-tekla-s>), Editor

Subscribe to RSS Feed (<https://spectrum.ieee.org/rss/blog/view-from-the-valley/fulltext>)